


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used: virus scan file path cluster

Found 89 of 207,474

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 89

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [On incremental file system development](#)


 Erez Zadok, Rakesh Iyer, Nikolai Joukov, Gopalan Sivathanu, Charles P. Wright
 May 2006 **ACM Transactions on Storage (TOS)**, Volume 2 Issue 2

Publisher: ACM Press

 Full text available: pdf(260.40 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Developing file systems from scratch is difficult and error prone. Using layered, or stackable, file systems is a powerful technique to incrementally extend the functionality of existing file systems on commodity OSes at runtime. In this article, we analyze the evolution of layering from historical models to what is found in four different present day commodity OSes: Solaris, FreeBSD, Linux, and Microsoft Windows. We classify layered file systems into five types based on their functionality and ...

Keywords: I/O manager, IRP, Layered file systems, VFS, extensibility, stackable file systems, vnode

2 [Document detection: TIPSTER phase I final report](#)

Bill Caid, Stephen Gallant, Joel Carleton, David Sudbeck

 September 1993 **Proceedings of a workshop on held at Fredericksburg, Virginia: September 19-23, 1993**

Publisher: Association for Computational Linguistics

 Full text available: pdf(1.84 MB) Additional Information: [full citation](#), [abstract](#)

During Phase I of the TIPSTER program, HNC developed a unique approach to machine learning of similarity of meaning. This approach, embodied in a system called "MatchPlus", exploits this learned similarity of meaning for concept-based text retrieval, routing and visualization of textual information. MatchPlus uses an information representation scheme called "context vectors" to encode similarity of usage. Key attributes of the context vector approach are as follows: • Words, documents, and q ...

3 [The internet worm program: an analysis](#)



Eugene H. Spafford

 January 1989 **ACM SIGCOMM Computer Communication Review**, Volume 19 Issue 1

Publisher: ACM Press

 Full text available: pdf(2.45 MB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

On the evening of 2 November 1988, someone infected the Internet with a *worm*

program. That program exploited flaws in utility programs in systems based on BSD-derived versions of UNIX. The flaws allowed the program to break into those machines and copy itself, thus *infecting* those systems. This program eventually spread to thousands of machines, and disrupted normal activities and Internet connectivity for many days. This report gives a detailed description of the components of the ...

4 Distributed Worm Simulation with a Realistic Internet Model

Songjie Wei, Jelena Mirkovic, Martin Swamy

June 2005 **Proceedings of the 19th Workshop on Principles of Advanced and Distributed Simulation PADS '05**

Publisher: IEEE Computer Society

Full text available:  [pdf\(326.02 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Internet worm spread is a phenomenon involving millions of hosts, who interact in complex and diverse environment. Scanning speed of each infected host depends on its resources and the defenses at work in its network. Aggressive worms further interact with the underlying Internet topology .. the dynamics of the spread is constrained by the limited bandwidth of network links, and high-volume scan traffic leads to BGP router failure thus affecting global routing. Worm traffic also interacts with I ...

5 Graph mining: Laws, generators, and algorithms



Deepayan Chakrabarti, Christos Faloutsos

June 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(910.68 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

How does the Web look? How could we tell an abnormal social network from a normal one? These and similar questions are important in many fields where the data can intuitively be cast as a graph; examples range from computer networks to sociology to biology and many more. Indeed, any $M : N$ relation in database terminology can be represented as a graph. A lot of these questions boil down to the following: "How can we generate synthetic but realistic graphs?" To answer thi ...

Keywords: Generators, graphs, patterns, social networks

6 Computer security and encryption II: Scanning workstation memory for malicious codes using dedicated coprocessors



Sirish A. Kondi, Yoginder S. Dandass

March 2006 **Proceedings of the 44th annual Southeast regional conference ACM-SE 44**

Publisher: ACM Press

Full text available:  [pdf\(176.91 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the implementation of a coprocessor platform for scanning workstation memory in order to detect signatures of malicious codes. The coprocessor is especially beneficial in clusters of workstations used for high performance computing where the overhead imposed by software-based intrusion detection codes is unacceptable. The coprocessor connects to the host via the PCI bus and accesses the host's memory using bus mastering DMA. The coprocessor interprets the host's virtual memor ...

Keywords: FPGA, coprocessor, intrusion detection, signature matching

7 A survey of peer-to-peer content distribution technologies

Stephanos Androutsellis-Theotokis, Diomidis Spinellis



December 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 4

Publisher: ACM Press

Full text available: [pdf\(517.77 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance. Content distribution is an important peer-to-peer application ...

Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

8 Traffic characterization: Characteristics of internet background radiation



Ruoming Pang, Vinod Yegneswaran, Paul Barford, Vern Paxson, Larry Peterson

October 2004 **Proceedings of the 4th ACM SIGCOMM conference on Internet measurement IMC '04**

Publisher: ACM Press

Full text available: [pdf\(396.12 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Monitoring any portion of the Internet address space reveals incessant activity. This holds even when monitoring traffic sent to unused addresses, which we term "background radiation." Background radiation reflects fundamentally nonproductive traffic, either malicious (flooding backscatter, scans for vulnerabilities, worms) or benign (misconfigurations). While the general presence of background radiation is well known to the network operator community, its nature has yet to be broadly charac ...

Keywords: honeypot, internet background radiation, network telescope

9 Implementing sorting in database systems



Goetz Graefe

September 2006 **ACM Computing Surveys (CSUR)**, Volume 38 Issue 3

Publisher: ACM Press

Full text available: [pdf\(518.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Most commercial database systems do (or should) exploit many sorting techniques that are publicly known, but not readily available in the research literature. These techniques improve both sort performance on modern computer systems and the ability to adapt gracefully to resource fluctuations in multiuser operations. This survey collects many of these techniques for easy reference by students, researchers, and product developers. It covers in-memory sorting, disk-based external sorting, and cons ...

Keywords: Key normalization, asynchronous read-ahead, compression, dynamic memory resource allocation, forecasting, graceful degradation, index operations, key conditioning, nested iteration

10 Improved error reporting for software that uses black-box components



Jungwoo Ha, Christopher J. Rossbach, Jason V. Davis, Indrajit Roy, Hany E. Ramadan, Donald E. Porter, David L. Chen, Emmett Witchel

June 2007 **ACM SIGPLAN Notices , Proceedings of the 2007 ACM SIGPLAN conference on Programming language design and implementation PLDI '07**, Volume 42 Issue 6

Publisher: ACM Press

Full text available:  [pdf\(345.48 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

An error occurs when software cannot complete a requested action as a result of some problem with its input, configuration, or environment. A high-quality error report allows a user to understand and correct the problem. Unfortunately, the quality of error reports has been decreasing as software becomes more complex and layered. End-users take the cryptic error messages given to them by programs and struggle to fix their problems using search engines and support websites. Developers cannot imp ...

Keywords: classification, error report, machine learning, profiling, software support

11 The effects of information scent on visual search in the hyperbolic tree browser

 Peter Pirolli, Stuart K. Card, Mija M. Van Der Wege

March 2003 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 10 Issue 1


Publisher: ACM Press

Full text available:  [pdf\(2.37 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Hyperbolic Tree is a focus + context information visualization that has been developed to amplify users' ability to navigate large tree-structured information systems. Information scent is a theoretical construct that captures one kind of interaction between task and display. Information scent is provided by task-relevant display cues, such as node labels on a tree that influence a user's visual search behavior and navigation decisions. An empirical Accuracy of Scent (AOS) score was developed ...


Keywords: Hyperbolic Tree, Information visualization, fisheye-lens visual search, focus+context, information foraging, information scent, interactive computer graphics

12 Astrolabe: A robust and scalable technology for distributed system monitoring, management, and data mining

 Robbert Van Renesse, Kenneth P. Birman, Werner Vogels

May 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 2


Publisher: ACM Press

Full text available:  [pdf\(341.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Scalable management and self-organizational capabilities are emerging as central requirements for a generation of large-scale, highly dynamic, distributed applications. We have developed an entirely new distributed information management system called Astrolabe. Astrolabe collects large-scale system state, permitting rapid updates and providing on-the-fly attribute aggregation. This latter capability permits an application to locate a resource, and also offers a scalable way to track sys ...

Keywords: Aggregation, epidemic protocols, failure detection, gossip, membership, publish-subscribe, scalability

13 The elements of nature: interactive and realistic techniques

 Oliver Deussen, David S. Ebert, Ron Fedkiw, F. Kenton Musgrave, Przemyslaw Prusinkiewicz, Doug Roble, Jos Stam, Jerry Tessendorf

August 2004 **ACM SIGGRAPH 2004 Course Notes SIGGRAPH '04**

Publisher: ACM Press

Full text available:  [pdf\(17.65 MB\)](#) Additional Information: [full citation](#), [abstract](#)

This updated course on simulating natural phenomena will cover the latest research and production techniques for simulating most of the elements of nature. The presenters will provide movie production, interactive simulation, and research perspectives on the difficult task of photorealistic modeling, rendering, and animation of natural phenomena. The course offers a nice balance of the latest interactive graphics hardware-based simulation techniques and the latest physics-based simulation techni ...

14 A holistic approach to service survivability



Angelos D. Keromytis, Janak Parekh, Philip N. Gross, Gail Kaiser, Vishal Misra, Jason Nieh, Dan Rubenstein, Sal Stolfo

October 2003 **Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems: in association with 10th ACM Conference on Computer and Communications Security SSRS '03**

Publisher: ACM Press

Full text available: [pdf\(1.58 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present SABER (Survivability Architecture: Block, Evade, React), a proposed survivability architecture that blocks, evades and reacts to a variety of attacks by using several security and survivability mechanisms in an automated and coordinated fashion. Contrary to the ad hoc manner in which contemporary survivable systems are built-using isolated, independent security mechanisms such as firewalls, intrusion detection systems and software sandboxes-SABER integrates several different techno ...

Keywords: intrusion detection, overlay networks, survivability

15 VizSEC state analysis session: NVisionIP: netflow visualizations of system state for security situational awareness



Kiran Lakkaraju, William Yurcik, Adam J. Lee

October 2004 **Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security VizSEC/DMSEC '04**

Publisher: ACM Press

Full text available: [pdf\(693.53 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The number of attacks against large computer systems is currently growing at a rapid pace. Despite the best efforts of security analysts, large organizations are having trouble keeping on top of the current state of their networks. In this paper, we describe a tool called NVisionIP that is designed to increase the security analyst's situational awareness. As humans are inherently visual beings, NVisionIP uses a graphical representation of a class-B network to allow analysts to quickly visuali ...

Keywords: NetFlows, security system state, security visualization, situational awareness

16 Data Mining and Predictive Modeling of Biomolecular Network from Biomedical Literature Databases

Xiaohua Hu, Daniel D. Wu

April 2007 **IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)**, Volume 4 Issue 2

Publisher: IEEE Computer Society Press

Full text available: [pdf\(3.81 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

In this paper, we present a novel approach Bio-IEDM (Biomedical Information Extraction and Data Mining) to integrate text mining and predictive modeling to analyze biomolecular network from biomedical literature databases. Our method consists of two phases. In

phase 1, we discuss a semisupervised efficient learning approach to automatically extract biological relationships such as protein-protein interaction, protein-gene interaction from the biomedical literature databases to construct the biom ...

Keywords: Biomolecular network, semisupervised learning, scale-free network, information extraction, biological complexes (communities).

17 Intrusion detection: Specification-based anomaly detection: a new approach for detecting network intrusions



R. Sekar, A. Gupta, J. Frullo, T. Shanbhag, A. Tiwari, H. Yang, S. Zhou
November 2002 **Proceedings of the 9th ACM conference on Computer and communications security CCS '02**

Publisher: ACM Press

Full text available: pdf(127.45 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Unlike signature or misuse based intrusion detection techniques, anomaly detection is capable of detecting novel attacks. However, the use of anomaly detection in practice is hampered by a high rate of false alarms. Specification-based techniques have been shown to produce a low rate of false alarms, but are not as effective as anomaly detection in detecting novel attacks, especially when it comes to network probing and denial-of-service attacks. This paper presents a new approach that combines ...

Keywords: anomaly detection, intrusion detection, network monitoring

18 1 - Regular Articles: Average-optimal single and multiple approximate string matching



Kimmo Fredriksson, Gonzalo Navarro
December 2004 **Journal of Experimental Algorithmics (JEA)**, Volume 9

Publisher: ACM Press

Full text available: pdf(1.77 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present a new algorithm for multiple approximate string matching. It is based on reading backwards enough l-grams from text windows so as to prove that no occurrence can contain the part of the window read, and then shifting the window. We show analytically that our algorithm is optimal on average. Hence our first contribution is to fill an important gap in the area, since no average-optimal algorithm existed for multiple approximate string matching. We consider several variants and practical i ...

Keywords: Algorithms, approximate string matching, biological sequences, multiple string matching, optimality

19 Risks to the public in computers and related systems



Peter G. Neumann
January 1990 **ACM SIGSOFT Software Engineering Notes**, Volume 15 Issue 1

Publisher: ACM Press


Full text available: pdf(2.11 MB) Additional Information: [full citation](#)

20 Formation and simulation: Worm anatomy and model



Dan Ellis
October 2003 **Proceedings of the 2003 ACM workshop on Rapid malware WORM '03**

Publisher: ACM Press

Full text available:  [pdf\(273.58 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present a general framework for reasoning about network worms and analyzing the potency of worms within a specific network. First, we present a discussion of the life cycle of a worm based on a survey of contemporary worms. We build on that life cycle by developing a relational model that associates worm parameters, attributes of the environment, and the subsequent potency of the worm. We then provide a worm analytic framework that captures the generalized mechanical process a worm goes through ...

Keywords: network modeling, network security, Turing machine, worm

Results 1 - 20 of 89

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used: virus scan cluster

Found 365 of 207,474

 Sort results
by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

 Display
results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Industry/government track poster: Short term performance forecasting in enterprise systems](#)



Rob Powers, Moises Goldszmidt, Ira Cohen

 August 2005 **Proceeding of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining KDD '05**

Publisher: ACM Press

 Full text available: [pdf\(882.75 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We use data mining and machine learning techniques to predict upcoming periods of high utilization or poor performance in enterprise systems. The abundant data available and complexity of these systems defies human characterization or static models and makes the task suitable for data mining techniques. We formulate the problem as one of classification: given current and past information about the system's behavior, can we forecast whether the system will meet its performance targets over the ne ...

Keywords: enterprise systems, performance forecasting

2 [Defensive techniques: Proactive security for mobile messaging networks](#)



Abhijit Bose, Kang G. Shin

 September 2006 **Proceedings of the 5th ACM workshop on Wireless security WiSe '06**

Publisher: ACM Press

 Full text available: [pdf\(281.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Interoperability of IM (Instant Messaging) and SMS (Short Messaging Service) networks allows users to seamlessly use a variety of computing devices from desktops to cellular phones and mobile handhelds. However, this increasing convergence has also attracted the attention of malicious software writers. In the past few years, the number of malicious codes that target messaging networks, primarily IM and SMS, has been increasing exponentially. Large message volume and number of users in these ...

Keywords: Instant Messaging (IM), SMS/MMS, containment, mobile viruses, proactive security, worms

3 [A high-availability high-performance e-mail cluster](#)



Wyman Miles

 November 2002 **Proceedings of the 30th annual ACM SIGUCCS conference on User**

services SIGUCCS '02**Publisher:** ACM PressFull text available:  [pdf\(175.10 KB\)](#) Additional Information: [full citation](#), [index terms](#)

Keywords: anti-spam, anti-virus, cluster, electronic mail, failover, high-availability, high-performance, mail routing, proxy, redundancy

4 Computer security and encryption II: Scanning workstation memory for malicious codes using dedicated coprocessors



Sirish A. Kondi, Yoginder S. Dandass

March 2006 **Proceedings of the 44th annual Southeast regional conference ACM-SE 44****Publisher:** ACM PressFull text available:  [pdf\(176.91 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the implementation of a coprocessor platform for scanning workstation memory in order to detect signatures of malicious codes. The coprocessor is especially beneficial in clusters of workstations used for high performance computing where the overhead imposed by software-based intrusion detection codes is unacceptable. The coprocessor connects to the host via the PCI bus and accesses the host's memory using bus mastering DMA. The coprocessor interprets the host's virtual memory ...

Keywords: FPGA, coprocessor, intrusion detection, signature matching

5 Network security: Code red worm propagation modeling and analysis



Cliff Changchun Zou, Weibo Gong, Don Towsley

November 2002 **Proceedings of the 9th ACM conference on Computer and communications security CCS '02****Publisher:** ACM PressFull text available:  [pdf\(197.17 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Code Red worm incident of July 2001 has stimulated activities to model and analyze Internet worm propagation. In this paper we provide a careful analysis of Code Red propagation by accounting for two factors: one is the dynamic countermeasures taken by ISPs and users; the other is the slowed down worm infection rate because Code Red rampant propagation caused congestion and troubles to some routers. Based on the classical epidemic Kermack-McKendrick model, we derive a general Internet worm model ...

Keywords: epidemic model, internet worm modeling, two-factor worm model

6 Surviving threats: Locality: a new paradigm for thinking about normal behavior and outsider threat



John McHugh, Carrie Gates

August 2003 **Proceedings of the 2003 workshop on New security paradigms NSPW '03****Publisher:** ACM PressFull text available:  [pdf\(760.75 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Locality as a unifying concept for understanding the normal behavior of benign users of computer systems is suggested as a unifying paradigm that will support the detection of malicious anomalous behaviors. The paper notes that locality appears in many dimensions and applies to such diverse mechanisms as the working set of IP addresses contacted

during a web browsing session, the set of email addresses with which one customarily corresponds, the way in which pages are fetched from a web site. In ...

Keywords: locality, network observation, system behavior

7 Synthesizing Realistic Computational Grids

Dong Lu, Peter A. Dinda

November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing SC '03**

Publisher: IEEE Computer Society

Full text available:  [pdf\(224.44 KB\)](#) Additional Information: [full citation](#), [abstract](#)

Realistic workloads are essential in evaluating middleware for computational grids. One important component is the raw grid itself: a network topology graph annotated with the hardware and software available on each node and link. This paper defines our requirements for grid generation and presents GridG, our extensible generator. We describe GridG in two steps: topology generation and annotation. For topology generation, we have both model and mechanism. We extend Tiers, an existing tool from t ...


8 Internet intrusions: global characteristics and prevalence



Vinod Yegneswaran, Paul Barford, Johannes Ullrich

June 2003 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '03**, Volume 31 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(699.44 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Network intrusions have been a fact of life in the Internet for many years. However, as is the case with many other types of Internet-wide phenomena, gaining insight into the *global* characteristics of intrusions is challenging. In this paper we address this problem by systematically analyzing a set of firewall logs collected over four months from over 1600 different networks world wide. The first part of our study is a general analysis focused on the issues of distribution, categorization ...

Keywords: internet performance and monitoring, network security, wide area measurement

9 Document detection: TIPSTER phase I final report

Bill Caid, Stephen Gallant, Joel Carleton, David Sudbeck

September 1993 **Proceedings of a workshop on held at Fredericksburg, Virginia: September 19-23, 1993**

Publisher: Association for Computational Linguistics

Full text available:  [pdf\(1.84 MB\)](#) Additional Information: [full citation](#), [abstract](#)

During Phase I of the TIPSTER program, HNC developed a unique approach to machine learning of similarity of meaning. This approach, embodied in a system called "MatchPlus", exploits this learned similarity of meaning for concept-based text retrieval, routing and visualization of textual information. MatchPlus uses an information representation scheme called "context vectors" to encode similarity of usage. Key attributes of the context vector approach are as follows: • Words, documents, and q ...

10 Web search 3: Improving web search results using affinity graph

Benyu Zhang, Hua Li, Yi Liu, Lei Ji, Wensi Xi, Weiguo Fan, Zheng Chen, Wei-Ying Ma



August 2005 **Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '05**

Publisher: ACM Press

Full text available: [pdf\(326.20 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper, we propose a novel ranking scheme named Affinity Ranking (AR) to re-rank search results by optimizing two metrics: (1) diversity -- which indicates the variance of topics in a group of documents; (2) information richness -- which measures the coverage of a single document to its topic. Both of the two metrics are calculated from a directed link graph named Affinity Graph (AG). AG models the structure of a group of documents based on the asymmetric content similarities between each ...

Keywords: affinity ranking, diversity, information retrieval, information richness, link analysis

11 E-commerce and e-content: Detectives: detecting coalition hit inflation attacks in advertising networks streams



Ahmed Metwally, Divyakant Agrawal, Amr El Abbadi

May 2007 **Proceedings of the 16th international conference on World Wide Web WWW '07**

Publisher: ACM Press

Full text available: [pdf\(292.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Click fraud is jeopardizing the industry of Internet advertising. Internet advertising is crucial for the thriving of the entire Internet, since it allows producers to advertise their products, and hence contributes to the well being of e-commerce. Moreover, advertising supports the intellectual value of the Internet by covering the running expenses of publishing content. Some content publishers are dishonest, and use automation to generate traffic to defraud the advertisers. Similarly, some ...

Keywords: approximate set similarity, click spam detection, cliques enumeration, coalition fraud attacks, real data experiments, similarity-sensitive sampling

12 CorMet: a computational, corpus-based conventional metaphor extraction system

Zachary J. Mason

March 2004 **Computational Linguistics**, Volume 30 Issue 1

Publisher: MIT Press

Full text available: [pdf\(246.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

CorMet is a corpus-based system for discovering metaphorical mappings between concepts. It does this by finding systematic variations in domain-specific selectional preferences, which are inferred from large, dynamically mined Internet corpora. Metaphors transfer structure from a source domain to a target domain, making some concepts in the target domain metaphorically equivalent to concepts in the source domain. The verbs that select for a concept in the source domain tend to select for its meta ...

13 What's Strange About Recent Events (WSARE): An Algorithm for the Early Detection of Disease Outbreaks

Weng-Keen Wong, Andrew Moore, Gregory Cooper, Michael Wagner

December 2005 **The Journal of Machine Learning Research**, Volume 6

Publisher: MIT Press

Full text available: [pdf\(341.72 KB\)](#) Additional Information: [full citation](#), [abstract](#)

Traditional biosurveillance algorithms detect disease outbreaks by looking for peaks in a

univariate time series of health-care data. Current health-care surveillance data, however, are no longer simply univariate data streams. Instead, a wealth of spatial, temporal, demographic and symptomatic information is available. We present an early disease outbreak detection algorithm called What's Strange About Recent Events (WSARE), which uses a multivariate approach to improve its timeliness of detect ...

14 Attacks and cryptanalysis: Puppetnets: misusing web browsers as a distributed attack infrastructure



V. T. Lam, S. Antonatos, P. Akritidis, K. G. Anagnostakis

October 2006 **Proceedings of the 13th ACM conference on Computer and communications security CCS '06**

Publisher: ACM Press

Full text available: pdf(871.35 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Most of the recent work on Web security focuses on preventing attacks that *directly* harm the browser's host machine and user. In this paper we attempt to quantify the threat of browsers being *indirectly* misused for attacking third parties. Specifically, we look at how the existing Web infrastructure (e.g., the languages, protocols, and security policies) can be exploited by malicious Web sites to remotely instruct browsers to orchestrate actions including denial of service attacks, ...

Keywords: distributed attacks, malicious software, web security

15 Towards a national infectious disease information infrastructure: a case study in West Nile virus and botulism

Daniel Zeng, Hsinchun Chen, Chunju Tseng, Catherine A. Larson, Millicent Eidson, Ivan Gotham, Cecil Lynch, Michael Ascher

May 2004 **Proceedings of the 2004 annual national conference on Digital government research dg.o '04**

Publisher: Digital Government Research Center

Full text available: pdf(284.11 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Information technologies are playing an increasingly important role in preventing, detecting, and managing infectious disease outbreaks. This paper presents a collaborative infectious disease informatics project led by an interdisciplinary team of information systems researchers and public health researchers and practitioners. This project has resulted in a research prototype called the WNV-BOT Portal system, which provides an integrated infectious disease information sharing, analysis, and visu ...

16 Distributed Worm Simulation with a Realistic Internet Model

Songjie Wei, Jelena Mirkovic, Martin Swamy

June 2005 **Proceedings of the 19th Workshop on Principles of Advanced and Distributed Simulation PADS '05**

Publisher: IEEE Computer Society


Full text available: pdf(326.02 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Internet worm spread is a phenomenon involving millions of hosts, who interact in complex and diverse environment. Scanning speed of each infected host depends on its resources and the defenses at work in its network. Aggressive worms further interact with the underlying Internet topology .. the dynamics of the spread is constrained by the limited bandwidth of network links, and high-volume scan traffic leads to BGP router failure thus affecting global routing. Worm traffic also interacts with I ...

17 VizSEC link analysis session: VisFlowConnect: netflow visualizations of link relationships for security situational awareness

-  Xiaoxin Yin, William Yurcik, Michael Treaster, Yifan Li, Kiran Lakkaraju
October 2004 **Proceedings of the 2004 ACM workshop on Visualization and data mining for computer security VizSEC/DMSEC '04**

Publisher: ACM Press

Full text available:  [pdf\(1.51 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present a visualization design to enhance the ability of an administrator to detect and investigate anomalous traffic between a local network and external domains. Central to the design is a parallel axes view which displays NetFlow records as links between two machines or domains while employing a variety of visual cues to assist the user. We describe several filtering options that can be employed to hide uninteresting or innocuous traffic such that the user can focus his or her attention ...


Keywords: link analysis, link relationships, netflows, parallel axes, parallel coordinates, security, security visualization, situational awareness

- 18 Computer security (SEC): Unsupervised learning techniques for an intrusion detection system 

Stefano Zanero, Sergio M. Savaresi

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing SAC '04**

Publisher: ACM Press

Full text available:  [pdf\(337.99 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


With the continuous evolution of the types of attacks against computer networks, traditional intrusion detection systems, based on pattern matching and static signatures, are increasingly limited by their need of an up-to-date and comprehensive knowledge base. Data mining techniques have been successfully applied in host-based intrusion detection. Applying data mining techniques on raw network data, however, is made difficult by the sheer size of the input; this is usually avoided by discarding ...

Keywords: K-means, anomaly detection, intrusion detection, principal direction divisive partitioning, quality of clusters, self-organizing maps, unsupervised clustering

- 19 Behavior-based modeling and its application to Email analysis 

Salvatore J. Stolfo, Shlomo HersHKop, Chia-Wei Hu, Wei-Jen Li, Olivier Nimeskern, Ke Wang
May 2006 **ACM Transactions on Internet Technology (TOIT)**, Volume 6 Issue 2


Publisher: ACM Press

Full text available:  [pdf\(1.25 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The Email Mining Toolkit (EMT) is a data mining system that computes *behavior profiles or models* of user email accounts. These models may be used for a multitude of tasks including forensic analyses and detection tasks of value to law enforcement and intelligence agencies, as well for as other typical tasks such as virus and spam detection. To demonstrate the power of the methods, we focus on the application of these models to detect the early onset of a viral propagation without "c ...

Keywords: Email virus propagations, anomaly detection, behavior profiling

- 20 System demonstrations and posters: Surveillance geoinformatics of hotspot detection and prioritization for monitoring, etiology, early warning and sustainable management 

G. P. Patil, K. Sham Bhat, S. W. Joshi

May 2007 **Proceedings of the 8th annual international conference on Digital**

government research: bridging disciplines & domains dg.o '07**Publisher:** Digital Government Research CenterFull text available:  [pdf\(42.16 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The live year NSF DGP project has been instrumental to conceptualize surveillance geoinformatics partnership among several interested cross-disciplinary scientists in academia, agencies, and private sector. A declared need is around for statistical geoinformatics and software infrastructure for spatial and spatiotemporal hotspot detection. Our efforts are driven by a wide variety of case studies of potential interest to federal agencies involving critical society issues, such as public health ...

Keywords: decision support for hotspot detection, early warning, geosurveillance statistics, hotspot detection, prioritization, space-time hotspots, surveillance geoinformatics partnership

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)